EXHIBIT 6

REDACTED

UNITED STATES DISTRICT COURT

EASTERN DISTRICT OF TEXAS

SHERMAN DIVISION

The State of Texas, et. al. Plaintiff, v. Google LLC, Defendant.	Case No: 4:20-cv-00957-SDJ

Expert Report of Matthew Weinberg

6/7/2024

Matthew Weinberg

in conducting my work and forming my opinions in this case. I reserve the right to supplement or amend this report if my opinions change or require supplementation as a result of my ongoing review of documents.

D. Summary of Opinions

11. I have analyzed each of the conducts undertaken by Google and assessed how they change the auction procedure and how these changes affect the auction outcomes.

12. I conclude that:[2]

a. Google's implementation of Dynamic Allocation led to higher win rate and higher revenue for AdX[3] as well as lower win rate and lower revenue for non-Google exchanges. Furthermore, Enhanced Dynamic Allocation led to an increase in win rate and increase in revenue for AdX and reduced the value of direct deals for advertisers. Reducing the value of direct deals for advertisers would decrease the revenue earned by publishers via direct deals.

b. Header bidding improves publisher outcomes relative to the waterfall approach (with or without Dynamic Allocation and Enhanced Dynamic Allocation) and it can generate higher revenue for publishers compared to Exchange Bidding.

c. Unified Pricing Rules likely lead to lower revenue for the publishers. It also can lead to better win rate and revenue for Google's ad exchange AdX as well as Google's ad buying tools and lower the win rate and revenue for rival exchanges and ad buying tools.

d. Under the Dynamic Revenue Sharing (DRS) conduct,

    i. Dynamic Revenue Sharing version 1 (DRSv1) increased AdX win rate and revenue and decreased non-AdX exchanges' win rates and revenues, compared to no DRS,

---

[2] These conclusions discuss the isolated impacts of the conducts, and throughout the report I also provide conclusion regarding the interactions between the conducts.
[3] AdX is Google's ad exchange.

ii. Dynamic Revenue Sharing version 2 (DRSv2), in comparison to both no DRS and DRSv1, decreased advertiser payoff,[4] increased AdX win rate and revenue, decreased non-AdX exchange's win rates and revenues, and may also have decreased publisher revenue.

iii. Truthful Dynamic Revenue Sharing increased AdX win rate and revenue and decreased non-AdX exchange's win rates compared to no DRS,

iv. Google concealed information that is vital to advertisers and important to publishers by concealing DRSv1 from them.

e. Projects Bernanke and Global Bernanke did not affect GDN[5] advertisers and could increase some publishers' revenues while decreasing others. However, these projects also led to a lower win rate for non-GDN ad buying tools and advertisers that used those ad buying tools. Furthermore, Project Bernanke and Global Bernanke led to an increased win rate for GDN buyers (without improving GDN advertisers' payoffs), which leads to an increased win rate and revenue for GDN.

f. Reserve Price Optimization leads to higher revenue for Google's ad exchange AdX, and lower payoff to advertisers. It could also lead to lower payoff for some publishers. The negative effects of Reserve Price Optimization to advertiser payoff, and possibly some publishers' revenues, is due to (a) Google's concealment of the conduct during its initial rollout, and (b) barriers to publishers effectively setting reserve prices to optimize their revenue even after Google announced the conduct.

E. Methodology

13. Throughout the report, I am going to apply the mathematical principles, results and insights that stem from the canonical auction theory and game theory literatures. These methods of auction analysis are commonly accepted by researchers and practitioners across many different fields such as economics, computer science, and mathematics. Furthermore, these tools are commonly accepted by researchers and practitioners for the analysis of the market at hand, online display ads.

[4] I use the term "advertiser payoff" to refer the difference between the advertiser's value for the impression and the amount they pay for the impression.
[5] GDN refers to Google Display Network, Google's ad buying tool for small advertisers. Its current name is Google Ads.

exceeds r, and so does not need to be increased). In this case, the maximum of r and the highest other bid is the highest other bid, and the payment is the highest other bid.

iii. If the highest bid falls below r, then no bids survive, and therefore no one gets the item or pays the auctioneer.

26. In a first-price auction, there are two relevant ranges for the reserve: (1) the highest bid exceeds the reserve, in which case the auction concludes identically as if there were no reserve or (2) the highest bid falls below the reserve, in which case the auction is essentially nullified, and the item stays with the seller. In a second-price auction, there are three relevant ranges for the reserve: (1) the second-highest bid exceeds the reserve, where the auction concludes identically as if there were no reserve or (2) the second-highest bid falls below the reserve, but the highest bid exceeds the reserve, so the highest bidder still wins, but pays the reserve which is greater than the second-highest bid or (3) the highest bid falls below the reserve, so the auction is essentially nullified, and the item stays with the seller.

27. The leading example can be modified with reserve prices. Imagine that there are five bidders, who submit bids of $1, $8, $3, $5, $2, respectively. In a first-price auction with a reserve price of $4, Bidder Two wins and pays $8. In a second-price auction with reserve of $4, Bidder Two wins and pays $5. That is, both auctions conclude exactly as if there were no reserve, because the reserve is smaller than the second-highest bid. If the auctioneer sets a reserve of $6, then in a first-price auction, Bidder Two wins and pays $8. The first-price auction concludes exactly as if there were no reserve, because the reserve is smaller than the highest bid. In a second-price auction with a reserve of $6, Bidder Two wins and pays $6. That is, Bidder Two still wins because they outbid the reserve. However, Bidder Two pays more because the reserve is treated as the second-highest bid. Observe also that $6 is Bidder Two's minimum bid to win, because $6 exceeds all other bids, Bidder Two will win if and only if they submit a bid above the reserve of $6. If the reserve is $10, then in both the first- and the second-price auction the item remains unsold because the reserve exceeds the highest bid.[25] Figure 2 below illustrates this example.

[25] I further analyze this numeric example in the Appendix C.

win at a price above their value (beyond this, subject to winning, bidders prefer to pay as low a price as possible) and (b) bidders will strategize while bidding in attempt to get preferred outcomes. The complexity of strategies depends on how their values are formed (*i.e.,* private versus interdependent), and the auction format.

47. In order to understand how auction formats affect bidder strategies, further terminology is needed. A sealed bid single-item auction is **truthful** if each bidder receives the best possible outcome (given the other bidders' bids) by submitting a bid equal to their own value.[43]

48. I now provide an example to illustrate the concept of truthfulness. Imagine a single-item auction with two bidders. Bidder One has value $5 and Bidder Two has value $10. Imagine further that the auctioneer has chosen a first-price auction, and Bidder One has submitted a bid of $3. What is the best possible outcome for Bidder Two in such a setting? It may be tempting to first claim that the best possible outcome for Bidder Two is to win the item and pay $0. However, there is nothing Bidder Two can do to make this happen since Bidder One has submitted a bid of $3, so the only outcomes available to Bidder Two are to lose the item (by submitting a bid less than $3), or to win the item at a price greater than $3 (by submitting a bid $b that is higher than $3).[44] In particular, winning the item and paying $10 is certainly not the best possible outcome, which is the resulting outcome should Bidder Two bid their value (a better outcome, for example, would be win and pay $3.01 by submitting a bid of $3.01). Therefore, the first-price auction is not truthful.[45] The incentive of a bidder to submit a bid that is lower than their value for the item is called "**bid shading**." The amount by which the bidders are incentivized to shade their bids is determined by their comparison of the risk of paying more when they submit a higher bid and the risk of losing the auction when they submit a lower bid. Figure 6 below illustrates the ideas in this paragraph.

[43] More formally, Bidder i cannot control what bids are submitted by the other bidders. But, no matter what those bids are, Bidder i can go through the thought process of "which bid b_i gets me the best possible outcome, given that the other bidders have bid $b_1, \ldots, b_{i-1}, b_{i+1}, \ldots, b_n$?" An auction is truthful if the answer is always "submitting b_i equal to my value gets me the best possible outcome."

[44] Depending on how ties are broken, perhaps it is possible to win the item at exactly $3.

[45] In particular, note that an auction is truthful if it is always best to submit a bid equal to the value. Because this example witnesses one case where it is not, this example in fact constitutes a proof that the first-price auction is not truthful.

Figure 6: In a first-price auction, the bidders are incentivized to submit bids lower than their values

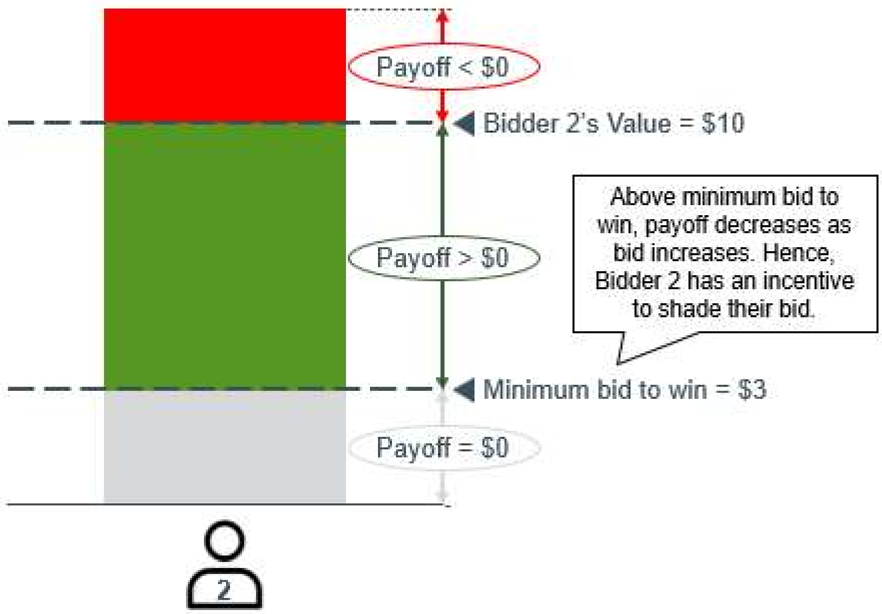

49. Imagine instead that the auctioneer has chosen a second-price auction, and Bidder One has submitted a bid of $3. In this case, the options available to Bidder Two are (a) lose (by submitting a bid less than $3), or (b) win and pay $3 (by submitting a bid greater than $3). Because Bidder Two's value is $10, their preferred outcome of the two possibilities is to win and pay $3. Submitting a bid of $10 is one way for Bidder Two to win and pay $3, and therefore bidding their value is one way for Bidder Two to get the best possible outcome.[46] The implications of these examples are not anomalies. Second-price auctions are truthful, and this is one of their key advantages in comparison to first-price auctions.[47, 48, 49]

[46] Still, recall that an auction is truthful if it is always best to submit a bid equal to the value. This example witnesses one case where it is, and so does not constitute a proof that the second-price auction is truthful. However, it does give intuition for why this is the case.
[47] A complete proof of this claim can be found in Theorem 1 the Appendix C.
[48] The truthfulness of the second-price auctions also holds with interdependent values, as long as the interdependency satisfies a technical condition called "single-crossing". Intuitively, the single-crossing condition assumes that Bidder i's value is more sensitive than Bidder j's value to Bidder i's signal (for example, car expert i places more emphasis on car expert i's inspection than car expert j places on car expert i's inspection). One precise mathematical statement is that for all signals $s_1,\ldots, s_n$, we have $\partial v_i(s_1,\ldots,s_n)/ \partial s_i \geq \partial v_j (s_1,\ldots,s_n)/ \partial s_i$. There are nuances to this generalization beyond the technical condition, it requires the auction designer to know how bidders map each other's signals to values and requires the auction designer to directly solicit signals rather than bids.
[49] Single-item auctions that charge minimum bids to win satisfy an even stronger property that bidding one's value is a dominant strategy, meaning there is a mathematically formal sense in which bidding the value is strictly better than any other bid. Formally, this just means that for any bid b not equal to a bidder's value: (a) bidding their value always gets the bidder an outcome at least as good as bidding b, no matter the other bids, and (b) there are some possible bids of the others where the bidder gets a strictly better outcome bidding their value than bidding b. The precise truthfulness properties that second-price auctions possess will not play a role in later analysis, this is stated primarily

1) Comparison of Incentives in First- and Second-Price Auctions

50. First- and second-price auctions differ in terms of the strategic behavior they elicit in the auction participants and the auctioneer. The suitable format depends on the context of the auction.

51. Given its truthfulness property, the second-price auction format may seem to be the ideal auction format compared to the first-price auctions. This is because bidding in auctions requires strategic sophistication. Bidders do not know other bidders' values, and so they do not know what bid will win them the item at the lowest possible price, in other words, their minimum bid to win. In a first-price auction, bidders would ideally want to bid a penny above their minimum bid to win (if their value exceeds that). However, in a truthful auction, the strategic sophistication is not needed because it is in the bidders' best interest to bid their own valuations of the auction item truthfully. Also, bidders tend to prefer straightforward auctions where they do not need to do any strategic bidding. This further enhances the benefit of the truthfulness of the second-price auction format.

52. However, there are potential mitigating factors as well. Truthfulness only holds when viewing this auction in isolation, it does not necessarily hold when considering a series of auctions for several reasons. One example to have in mind is when the auctioneer might change their reserve in later auctions based on bids in earlier auctions. Once the auctioneer's reserve is fixed, a second-price auction with that reserve is truthful, and a bidder's optimal bid for this auction is their value. However, if the bidder has a high value for the item and bids as such, this may indicate to the auctioneer that they should set a higher reserve in future auctions, and this higher reserve will certainly hurt the bidder in the future.[50] It may also signal to other bidders that similar items are valuable and increase their future values (and therefore bids).[51]

---

to note that the argument that bidders should bid their value in a truthful auction is slightly stronger than what is implied by Theorem 1 in the Appendix C.

[50] I believe this concept is relevant when analyzing Google's conduct called "Reserve Price Optimization." This is discussed in detail in Section IX.

[51] Additionally, Myerson's (1979) seminal work introduces the so-called "revelation principle", which essentially provides a truthful wrapper to put around a non-truthful auction to handle the non-truthfulness on behalf of the bidder. Essentially, imagine that the real bidder hires a sophisticated data-rich consultant to bid in a non-truthful auction on their behalf, tells the consultant their value, and lets the consultant take care of optimizing the bid. From the bidder's perspective, assuming the consultant is honest and effective, this is now just a truthful auction (the bidder is best off giving the consultant the most accurate information to work with and letting them do their job), and only the consultant deals with the messy optimization. In settings where such a sophisticated data-rich consultant is widely available, the additional benefit of truthfulness is smaller. Roger B. Myerson. "Incentive Compatibility and the Bargaining Problem." *Econometrica* vol. 47, no. 1. 1979. pg. 61–73.

58. Information also helps bidders to optimize their bids. One benefit of a truthful auction is that the only "hard part" of the auction is determining the value. Once the value is determined, it is optimal to bid that value. But many auctions are non-truthful. In a non-truthful auction, optimal bidding is challenging since, by definition, it requires knowing competitors' bids. For example, imagine a bidder in a first-price auction, and their value is $5. What should they bid? They definitely should not bid anything above $5. But other than that, any strategy for this bidder depends on other bidders' bids. But, if somehow the others' bids can be seen, then bid optimization becomes a straightforward process (either bidding a penny more than the maximum bid, or purposefully losing if the maximum bid exceeds the willingness to pay).[54] Hence, the ability to see others' bids, especially if the other bidders cannot see rival bids, would be a great strategic advantage.[55] Similarly, there would be benefit from any data that helps predict other bidders' bids (perhaps, as an example, their historical bids submitted on similar items).

59. Lastly, the sellers use information when they are determining appropriate reserve prices for their items. For example, in the context of a second-price auction, a too high reserve will nullify the entire auction, a too low reserve has no impact, and setting a reserve in the sweet spot between the highest and second-highest bid yields extra revenue. For example, if the highest bid is $20, and the second-highest bid is $15, the seller would ideally like to set a reserve at exactly $20. Failing that, they would really like to avoid setting a reserve above $20 and prefer to set larger reserves between $15 and $20. But if instead the highest bid was $100, and the second-highest bid was $50, the seller would ideally like to set a reserve exactly at $100, definitely not over $100, and somewhere between $50 and $100. Notice that a good reserve in the first case ($20) is useless in the second, while a good reserve (greater than $50) in the second case nullifies all revenue in the first. *A priori*, with no further information, deciding on the optimal reserve is challenging at the very least. But every time a seller sees bids in an auction for a similar item, they learn a little bit about what they might expect the next time. This data is valuable, because it allows the seller to predict whether they are more likely to be shooting for a reserve between $15 and $20 or between $50 and $100, and to target their reserve at the likely case.

60. Another crucial concept is **information asymmetry**, which refers to cases where one agent has more information than another. One source of information asymmetry might be if one special bidder gets to see others' bids in a first-price auction before submitting their own. The

[54] In this example, if the highest other bid is $3.38, they should bid $3.39. If the highest other bid is $1.28, they should bid $1.29.
[55] This point is relevant for "Last Look" combined with "Header Bidding" in Google's Ad Server DFP. These conducts will be discussed in detail later.

special bidder is better able to optimize their own bid, and also may be more informed about their own value if values are interdependent. Another example might be if one seller gets feedback from lots of auctions, but another does not. Then, the first seller is better informed and can potentially set better reserves. Imagine someone, while cleaning out their basement, finds their old collection of Pokémon cards and decides to sell them. They would like to set a reserve to make sure they optimize their revenue, but they might have no idea for how much Pokémon cards go these days. Any information gathered will be helpful to stop them from setting a reserve of $100 and risk getting fleeced, or a reserve of $100,000 and risk their collection going unsold. After setting a good reserve, they might launch a first-price auction. A friend of the seller might reach out and express their interest, but also the friend might be unable to both determine what their true value is (part of their value derives from their ability to further resell the set when they lose interest) and to bid strategically once a value is determined. In turn, the seller might decide that they will just share bids with the friend as they roll in and let them submit theirs at the end. This allows the friend to both (a) accurately form their value for the set, by observing the bids of others, and (b) bid optimally, by submitting a bid just a penny above the highest other bid (if they decide they want to win). By sharing this information with their friend, the seller has created an information asymmetry since the friend is both better informed about their value than other bidders, and better able to optimize their bidding strategy.

III. ONLINE DISPLAY ADVERTISING

61. In this section, I provide an overview of parties and concepts that arise within online display advertising.

A. Pertinent Products and Parties in Transactions for Online Display Ads

62. There are several players involved in online display advertising, the goal of this section is to outline those players as well as their goals and incentives.

1) Publishers and Ad Servers

63. **Publishers** are entities (e.g. the New York Times) that have webpages that can display ads and are therefore sellers of inventory, which is effectively the "eyeballs" of their users. Publishers (and third parties) have a range of information about the particular users visiting their page, and each unique visitor can be thought of as its own item for sale. The item for sale is referred to as an **impression**.

64. Like any other business selling goods, in the context of online display advertising, a publisher's primary goal is to make as much revenue as possible by selling their inventory. However, publishers must be mindful that their true goal is long-term revenue, and that sometimes actions that increase immediate revenue may be harmful to long-term revenue. Hence, publishers have concerns such as the quality of ads displayed in addition to immediate revenue. For example, if users find display ads offensive or annoying, they may negatively impact the publisher's brand and future revenue streams.

65. There are three key elements of the display ad sale process from the publisher's side. First, the primary benefit of the display ad ecosystem is that ads are targeted, which means that each user's eyeballs are treated as a truly distinct item. For example, the right to display an ad for running shoes to a runner is more valuable than showing that ad to a golfer. Second, the decision regarding which ad to display happens nearly instantaneously. Once the publisher learns of a user visiting its website, it must decide which ad to display by the time the webpage loads. Third, the publisher may not necessarily have a natural network of advertisers ready and waiting to bid on its impressions. The publisher must somehow reach interested advertisers.

66. Publishers face a challenging task when they are trying to sell impressions. Imagine the selling process of a publisher, who needs to track, store, and share targeting data on each user visiting their website, reach a wide network of potential advertisers, decide a revenue-maximizing auction to run, and do it all nearly instantaneously while the webpage loads. Publishers usually outsource these tasks to a dedicated product called an ad server.

67. An **ad server** is a service that helps publishers manage and sell inventory, overcoming the technical challenges listed above. Inventory management and optimal pricing are also data-intensive (to accurately determine the aggregate market for each impression) and mathematically sophisticated (to determine what the optimal auction is given the data).

68. An ad server's revenue typically comes either by charging publishers a fixed monthly rate, or charging a fee based on the volume of impressions served.[56] This suggests, in principle, that an ad server's primary goal is to serve their publisher's goals, since the ad server's monetary

[56] *See* AdGlare. "Plans & Pricing." Accessed on May 31, 2024. https://web.archive.org/web/20231203094651/https:/www.adglare.com/pricing (describing that AdGlare charges a monthly rate that is based on the number of ad requests.)

incentives are aligned with the publisher's monetary incentives.[57] Google operates an ad server called DoubleClick for Publishers (DFP), which was later merged into Google Ad Manager.[58]

69. There are two key ways through which an ad server might sell inventory. In a **live ad auction**, the ad server learns that a particular user is visiting the publisher's webpage, and while the page loads, runs an auction for the right to display an ad to this particular user. The auction begins only after the user is known, so potential advertisers can submit a bid based on the fine-grained information they learn about the user. In a **direct deal**, the publisher pre-arranges a contract with an advertiser to display their ad some number of times across some period at some predetermined price per impression, perhaps to users that satisfy some coarse targeting criteria. The ad server manages that deal as users visit the publisher's webpage. Because of this, while targeting criteria can still be used, it is coarser in comparison to the real-time data available in a live ad auction.[59, 60]

70. Establishing a relationship with advertisers for direct deals is time consuming and requires a business network. It is perhaps worth the effort for large publishers such as the New York Times and large advertisers such as Nike, but it is unlikely that smaller publishers (for example, a local food blog) grab the attention of Nike's marketing department for direct deals or smaller advertisers (for example, a local escape room) grab the attention of The New York Times for direct deals. Because live ad auctions are automated, there is no reason why code written by a local food blog cannot interact with code written by Nike.[61] As a result, live ad auctions enable small publishers

[57] An ad server that is not standalone (i.e., owned by an entity that operates elsewhere in the online display ad auction ecosystem) could certainly have alternate primary goals. These could be benign and still aligned with a publisher's if that entity views ad servers as a necessary part of the ecosystem for their primary service to thrive. But this also raises potential for misaligned incentives.

[58] Along with Google's own ad exchange, in 2018. This will be explained further below. *See* Jonathan Bellack. "Introducing Google Ad Manager" (June 27, 2018). Accessed on May 31, 2024. https://web.archive.org/web/20240112234145/https:/blog.google/products/admanager/introducing-google-ad-manager/

[59] NT Technology. "Why Is Targeting in Programmatic Ads Better Than Usual?" Accessed on June 5, 2024. https://web.archive.org/web/20240228203217/https://nt.technology/en/faq-en/why-is-targeting-in-programmatic-ads-better-than-usual/ ("Improved targeting capabilities. Through programmatic advertising, you can easily target the specific audience you want to reach using all target opportunities. Rather than trying to reach sports car fans on an auto site, brands have the opportunity to create an audience segment of sports car fans and reach them across hundreds of websites, wherever they happen to be online.")

[60] In reality, there are more types of trade in online display advertising markets, but these two are the most relevant ones to the case, so I choose to focus on these. *See generally* Google. "Line item types and priorities." Accessed on May 31, 2024. https://web.archive.org/web/20240216154938/https:/support.google.com/admanager/answer/177279?hl=en

[61] But even this is still not trivial, since something still has to do the work of finding and connecting these two codebases together (and quickly, by the time a webpage loads). This challenge motivates the role of exchanges, which I discuss after introducing the buy side of the market.

to display ads from a wide range of advertisers without investing in the business network aspect of advertising.

71. A publisher can consider both a direct deal and a live ad auction for the same impression. For example, upon learning that a user is visiting their webpage, a publisher could first check if the impression satisfies a high-value direct deal and if so, sell it via direct deal. If not, the publisher could sell the same impression via live auction instead.

2) Advertisers, Large and Small Ad Buying Tools

72. **Advertisers** are entities that wish to display ads to users and are therefore buyers of inventory. Each advertiser has a distinct value[62] for each impression and that value is determined based on the information the advertiser learns about the user behind that impression. An advertiser's primary goal is to win impressions at a price below their value and to pay as low a price as possible.

73. Advertisers can purchase impressions through direct deals or live ad auctions, which can be challenging tasks. Imagine that Nike is looking to purchase targeted online advertising. They might engage in traditional marketing with large publishers via direct deals, but they will need to engage with live ad auctions to reach small and midsized publishers. Nike needs to process data on each impression to determine its value, manage an advertising budget across an extended time horizon, optimize bidding strategies in live ad auctions, and get connected to publishers in the first place. As a result, advertisers outsource these tasks to dedicated products called ad buying tools.

74. An **ad buying tool** is a service that helps advertisers find impressions that are available for sale, bid appropriately to balance the likelihood of winning versus price paid, manage a budget across a time horizon, process any available data on the impression to inform their value, and generally manage the process of purchasing impressions via live ad auctions. An ad buying tool's

[62] In the terminology introduced in Section II, I assume for the majority of this report that the advertisers have independent private values for impressions (the only exception is when identifying a potential impact of Enhanced Dynamic Allocation on direct deals in Section IV). This is a simplifying assumption (it is likely that no auction in real life purely abides by the independent private values model) that makes the analysis more tractable, and it is a sensible assumption to make since (a) internal Google documents demonstrate that Google assumes this as well (*e.g.*, GOOG-AT-MDL-004016180), (b) bidders are heterogeneous in their valuation for impressions (the impression from an avid runner is valued differently by Nike and McDonald's) and since the bidders do not know each other's identities, even if they learn about others' bids it could possibly not be that useful towards determining their own valuation. GOOG-AT-MDL-004016180 at -94. February 20, 2020. "Auction Theory Primer." ("Unlike with 2nd-price, in order to bid [in a first-price auction], *buyers must believe something about the competition!* **Assumption:** [emphasis in original] Independent Private Values model.")

primary goal is to optimize its revenue, which is typically earned as a fraction of payments made by the advertisers it serves (see the examples in paragraph 80 for a numerical illustration). Google has ad buying tools in two markets. Google Ads serves small advertisers and DV360 serves large advertisers.

75. Arguably the most important function of an ad buying tool is that it determines the bids on behalf of the advertisers, according to the goals they input to the system.[63] Since online ad auctions happen almost instantaneously, the advertisers themselves cannot possibly submit bids into the auctions. Instead, they input their goals for their advertisement campaigns into the ad buying tool, which then comes up with bids in a timely manner when impressions become available. The advertiser goals usually include parameters like the desired volume of impressions, budget and time horizon allocated for the campaign, and targeting criteria.[64]

3) Ad Exchanges

76. Publishers (via ad servers) and advertisers (via ad buying tools) form the sell side and buy side of the markets for live ad auctions. It is not a trivial process for ad servers and ad buying tools to find each other and transact. Even for something that is commonly bought and sold, such as a designer coat, finding every interested buyer on the internet is a difficult task. As a buyer, it is also a difficult task to scour the internet to find all the designer coats you are interested in. Hence a third-party market/exchange/bazaar would be relied on to aggregate supply and demand. For example, customers go to platforms like eBay for Pokémon cards, Etsy for engraved chopsticks, and Amazon for books. In all of these cases, the customers rely on the platform primarily to match them to sellers.[65] The market for impressions is no different, and ad exchanges exist to help publishers meet advertisers.

77. **Ad exchanges** provide the service of matching advertisers (buyers) to publishers (sellers). Ad servers contact an ad exchange with inventory for sale, and the exchange then connects to ad buying tools. Importantly, ad exchanges do not merely connect advertisers to publishers, they also run an auction to determine which advertiser wins the impression. That is, an ad exchange is more like the New York Stock Exchange (which specifies the stock trading mechanism) or Uber and Lyft (which specify the price at which riders and drivers transact) than a

[63] Depending on the advertisers they serve, the ad buying tools can enable varying degrees of advertiser input into the bidding algorithm.
[64] *See* Google. "Determine a bid strategy based on your goals." Accessed on June 6, 2024. https://web.archive.org/web/20240602100502/https://support.google.com/google-ads/answer/2472725
[65] Some platforms also offer derivative services, such as shipping, fraud protection, *etc.*

bazaar (which largely serves as a meeting point for buyers and sellers to engage in whatever interaction they like) or Craigslist (which functions largely, although not entirely, like a digital bazaar).

78. An ad exchange's primary goal is to optimize its revenue, which is typically earned as a fraction of payments made when it facilitates a transaction between an advertiser and publisher.[66] Google operates an exchange called AdX, which was later combined with DFP into Google Ad Manager.[67]

79. In sum, publishers sell impressions, and use ad servers to manage this process. Advertisers buy impressions and use ad buying tools to manage this process. Exchanges intermediate this process by connecting publishers to advertisers. Figure 7 below presents the participants and the intermediaries in the online display ads market, with each Google intermediary.

Figure 7: Participants and intermediaries in the online display ads market

80. In order to better understand each key player's role in the markets for the display ads, consider the following example. The New York Times is a publisher that uses DFP as its ad server

[66] The conducts described in this report allege examples where AdX claims to earn revenue as a fraction of payments made when it facilitates a purchase, but actually collects revenue through a more complicated mechanism.
[67] *See* Jonathan Bellack. "Introducing Google Ad Manager" (June 27, 2018). Accessed on May 31, 2024. https://web.archive.org/web/20240303134019/https:/www.blog.google/products/admanager/introducing-google-ad-manager/

Figure 8: Nike gets the impression by winning the AdX auction through its ad buying tool DV360

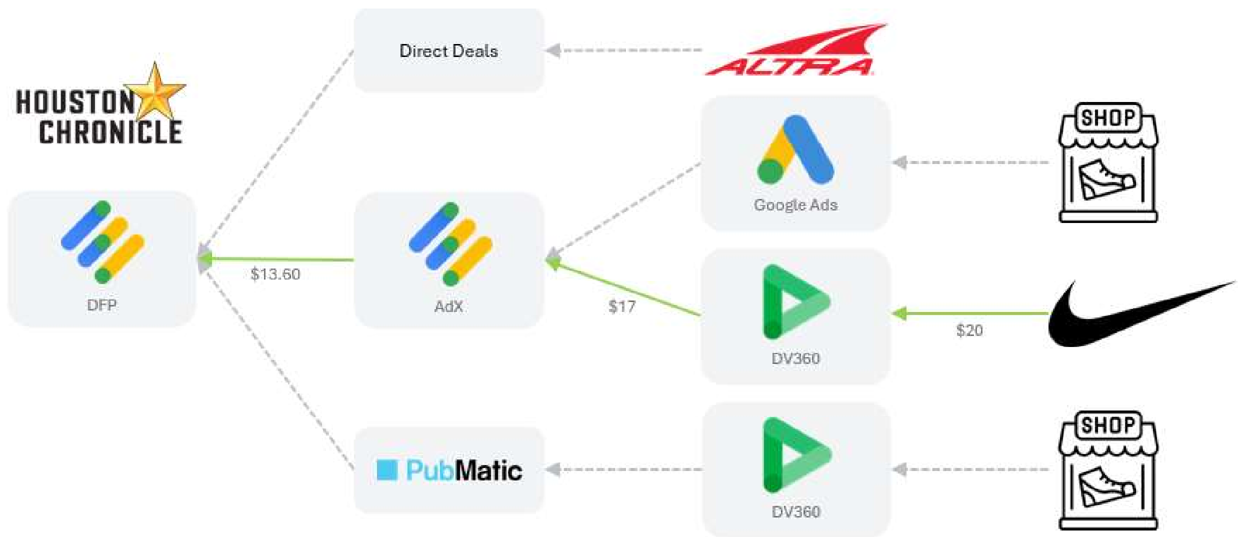

B. Auction Formats Unique to Online Display Advertising

81. In Section II, I overviewed standard auction concepts such as the first- and second-price auctions, reserves, and personalized reserves. These concepts are all directly relevant to the case at hand. This subsection overviews two further concepts that are central to the case and prevalent in the online display advertising ecosystem: the waterfall and header bidding.

82. One building block for both is the concept of **line items**. Within an ad server, a line item refers to a potential demand source (e.g. an advertiser or an intermediary for an advertiser). Line items can be complicated and contain many different types of information, such as (a) an ad to display in case this line item is selected, (b) information about the buyer and where to look for payment in case this line item is selected, (c) the price that would be paid in case this line item is selected, (d) criteria that determine which impressions are permitted to select this line item or (e) other metadata to aid the process.[71]

83. The following are a few examples of how demand sources exist as line items.

a. A **guaranteed direct deal** from an advertiser offers price *p* per impression for any impression that satisfies proposed coarse targeting criteria. The advertiser expects exactly some number of impressions to be displayed per time period. When a

[71] *See* Google. "About line items." Accessed on May 31, 2024. https://web.archive.org/web/20240216155011/https:/support.google.com/admanager/answer/9405477?hl=en (Google Ad Manager's online documentation on line items.)

publisher creates a guaranteed direct deal as a line item, it includes the targeting criteria, the price per impression, and where to look for payment. The ad server tracks the numbers to ensure that the right number of impressions are sold.[72]

b. A **non-guaranteed direct deal** from an advertiser offers price *p* per impression for any impression that satisfies proposed coarse targeting criteria. The advertiser might also place a cap on the number of times this direct deal can be fulfilled per time period. When a publisher creates a non-guaranteed direct deal as a line item, that line item includes the targeting criteria, price per impression, and where to look for payment. The ad server tracks to ensure that the cap is not exceeded.[73]

c. An **individual exchange** is also a line item. When a publisher wants to elicit bids from a specific exchange, they add the exchange as a line item. Note that there is a distinct line item for each exchange.[74]

1) The Waterfall

84. The first key auction concept unique to the online display advertising ecosystem is the **waterfall**, which is a process used by an ad server to sell an impression. When selling an impression via the waterfall, the ad server visits line items one at a time in a priority order set by the ad server and the publisher.[75] When a line item is selected as the winner of the auction, the waterfall concludes.

85. Ad servers typically prioritize direct deals ahead of ad exchanges. They first check whether the impression meets the coarse targeting criteria for a direct deal, and if so allocate the

[72] Guaranteed direct deals would be a type of what is called a "sponsorship line item." *See* Google. "Sponsorship line items." Accessed on May 31, 2024.
https://web.archive.org/web/20221209041446/https:/support.google.com/admanager/answer/177426?hl=en
Some guaranteed direct deals can be based on a percentage of all the impressions satisfying the targeting criteria as well, such as "25% of impressions coming from women in Plano, TX."
[73] Non-guaranteed direct deals would be a type of what is called a "price priority line item" in Google Ad Manager's online documentation. *See* Google. "Price Priority line items." Accessed on May 31, 2024.
https://web.archive.org/web/20240216154933/https:/support.google.com/admanager/answer/79306?hl=en
[74] These are called "ad exchange line items" in Google Ad Manager documentation. *See* Google. "Ad Exchange line items." Accessed on May 31, 2024.
https://web.archive.org/web/20221012075051/https:/support.google.com/admanager/answer/188523?hl=en
[75] The ad server determines the general groups of line items and the ranking among these groups. *See* Google. "Line item types and priorities." Accessed on May 31, 2024.
https://web.archive.org/web/20240216154938/https:/support.google.com/admanager/answer/177279?hl=en
Within these groups. the publisher might manually set a line item order or rely on the ad server for automatic ordering based on some metric, such as the value CPM. *See* Google. "Value CPM." Accessed on May 31, 2024.
https://web.archive.org/web/20221202071803/https:/support.google.com/admanager/answer/177222?hl=en

impression to the highest paying direct deal without visiting any exchanges.[76] If the impression is not selected for any direct deal, the ad server then visits exchanges one at a time, in an order set by the publisher, with a price floor of r set by the publisher, possibly different for different exchanges.[77] When visited, the exchange can either claim the impression (and pay at least r to the publisher) or pass.[78] Figure 9 below illustrates such an example. The waterfall format was in use by DFP since before DFP was purchased by Google.[79] Since the rise of header bidding in the 2010s, which is discussed below, the waterfall has played a relatively smaller role than it once played.[80]

---

[76] Prioritization of direct deals over exchanges is a curious feature of the waterfall. If I were to design a waterfall-like format from scratch and I were unconstrained by technological challenges, I would (a) find the maximum payment v I could get from a direct deal for this impression (maybe $v = 0$, if it satisfies no direct deal targeting criteria), then (b) visit exchanges in the waterfall but setting reserves informed by v (for example, I would certainly never set a reserve lower than v, because I would rather just fulfill a direct deal for price v than sell to an exchange below v and I might certainly set a reserve above v, in order to get a chance at even greater revenue) and then (c) if the waterfall completed without any exchange paying, I would use the impression to fulfill the direct deal and collect my v for doing so. I would do this because if I decide to take a direct deal without visiting exchanges, I limit myself to exactly v, while if I instead decide to visit exchanges with all reserves higher than v before deciding on the direct deal, I guarantee myself at least v (because I can always fall back on the direct deal if all exchanges pass), but have a shot at more than v (if I get lucky and an exchange bids). Similarly, if I decide to pass on a direct deal without visiting exchanges, I may wind up with 0 if all exchanges pass, while if I instead visit exchanges first and they all pass, I can still get something via direct deal. The documentation I have access to does not explicitly state why the opposite decision was made, but surrounding context clues suggest this was likely due to technical limitations in a novel ecosystem using software initially developed for a simpler ecosystem. Anecdotal evidence also suggests that direct deals typically paid much more per-impression than live ad auctions anyway, and so therefore the loss due to suboptimal ordering may have been minimal. *See* Google. "Understand Direct and Programmatic Ad Revenue." Accessed on May 31, 2024.
https://web.archive.org/web/20231226200704/https:/newsinitiative.withgoogle.com/resources/trainings/grow-digital-ad-revenue/understand-direct-and-programmatic-ad-revenue/
There could certainly be other reasons too. Still, I do not mean to imply that this curiosity has a simple resolution. Because direct deals are made with coarser targeting criteria than live ad auctions, direct deal advertisers may insist on being considered first in order to avoid becoming victims of "cream-skimming" (where among impressions that satisfy the same coarse targeting criteria, live ad auctions win the "good" impressions and leave the "bad" ones for direct deals). I share this commentary primarily to give an example of how an auction theorist might reason through the process of auction design, although this particular thought process plays a role in my later analysis of Enhanced Dynamic Allocation.

[77] These reserves can be set by the publisher on their ad server, or on the ad exchange integrated into their ad servers. Also, there are third party tools who provide revenue optimization services, in some cases they may set the reserve prices on behalf of the publishers. Throughout the report, I mostly abstract away from these differences, since they are functionally the same.

[78] Payments do not occur instantaneously on a per-impression basis. Publishers get periodical payouts depending on the ad server they choose to work with. *See* Google. "Ad Manager payment timelines." Accessed on May 31, 2024. https://web.archive.org/web/20221209003032/https:/support.google.com/admanager/answer/2671030?hl=en (Google's ad server payout options.)

[79] GOOG-NE-10780865 at -78, 79. May 5, 2020. "Clearing Up Misconceptions About Google's Ad Tech Business."

[80] *See* Sarah Sluis. "The Rise Of 'Header Bidding' And The End Of The Publisher Waterfall" (June 18, 2015). Accessed on May 31, 2024.
https://web.archive.org/web/20240216154913/https:/www.adexchanger.com/publishers/the-rise-of-header-bidding-and-the-end-of-the-publisher-waterfall/

89. DFP gives publishers the option to set whatever reserves they like.[83] Documentation suggests that most publishers chose to set reserves based on the historical average payment from that exchange.[84] Note that, given access to the distribution of historical payments, it is suboptimal to simply set the average payment as the reserve. This holds even when there is just a single exchange. Seminal work of Myerson (1981)[85] describes the revenue-maximizing reserve in this case. With multiple exchanges visited through the waterfall, optimal reserves are even more complex.

90. Documentation also suggests that publishers sort exchanges primarily in decreasing order of historical average payment, although exchange fill rate (the fraction of offers an exchange accepts) and ad quality (the quality of the visual displayed in the auctioned ad space) might play a role too.[86] This heuristic of setting the waterfall ordering on the basis of historical CPM[87] averages generates nontrivial incentives for the exchanges. A higher winning bid by exchanges increases their future reserves (because their average bid goes up), making future impressions more expensive. On the other hand, a higher reserve puts the exchange earlier in the waterfall, which gives them access to more impressions. Sorting exchanges in decreasing order of reserve is natural since (a) combined with the simple average CPM heuristic for setting reserves, exchanges have the indirect opportunity to pay more to be placed earlier in the waterfall, and (b) if instead exchanges eschew this opportunity and simply pay exactly the reserve to prevent the

---

[83] They were later constrained in their freedom to choose appropriate reserve prices by Unified Pricing Rules conduct, which I discuss in Section VI.
[84] GOOG-NE-10780865 at -81. May 5, 2020. "Clearing Up Misconceptions About Google's Ad Tech Business." ("Publishers typically set the net value CPM for their booked static remnant line items based on their estimates of what CPM the line item would likely generate (taking into account its historical performance) or based on a fixed-price the publisher had negotiated with a particular remnant demand partner.")
[85] Roger B. Myerson. "Optimal Auction Design." *Mathematics Of Operations Research* vol. 6, no. 1. 1981. pg. 58-73.
[86] *See* GOOG-NE-10780865 at -79. May 5, 2020. "Clearing Up Misconceptions About Google's Ad Tech Business." (¶3 describes how publishers ordered demand sources.)
[87] CPM refers to "cost per mille." It is the cost of purchasing one thousand impressions for an advertiser.

reserve price from increasing in future, the revenue-optimal ordering for the publishers is indeed to sort exchanges in decreasing order of reserve.[88, 89]

91. From an auction theory perspective, waterfalling as a procedure for selling impressions is inefficient and sacrifices revenue. This is because waterfalling forces the ad server to decide whether to sell to one exchange before learning what other exchanges might bid. A first- or second-price auction among exchanges (as opposed to the waterfalling) simultaneously considers all exchanges, which avoids this problem. Indeed, there are several research papers quantifying the suboptimality of sequential auction formats (specifically, posted-price mechanisms,[90] which are the closest common auction format to waterfalling) as compared to simultaneous auction formats (specifically, second-price auctions).[91] Moreover, the concept of sequential decision-making with partial information versus simultaneous decision-making with all information is well-studied within computer science broadly under the field of Online Algorithms (here, "online" refers to "making binding decisions one at a time with incomplete information" rather than "on the internet"),[92] and it is well-understood that binding sequential decisions come at a loss compared to a single decision with all the available information.

92. An ad server might still employ the waterfall process even with these suboptimalities. One reason computer scientists study online algorithms is because making a single decision with full

[88] This leads to the question of why exchanges ever pay more than the reserve price. First, some exchanges may have inflexible contracts with advertisers that preclude them from being particularly strategic with bids. For example, perhaps an exchange agrees to take exactly a 20% cut of the winning bid and give the rest to the publisher. Then, if this exchange has an advertiser willing to pay up to $6, and their reserve is $4, the highest revenue the exchange can collect is by paying $4.80 to the ad server and collecting $1.20 on $6. If instead the exchange were to pay exactly $4, this would correspond to taking a 20% cut of $5, which is just $1. However, this reasoning does not apply if exchanges were not bound by such contracts. An unbound exchange in this example could pay $4 to the ad server, collect $6 from the advertiser, and pocket $2. Second, if the ad server indeed uses the average of past clearing prices as a heuristic to set reserves, exchanges' incentives are complicated. If other exchanges are likely to purchase the impression, being early in the waterfall is the only way to get a shot at the impression. Therefore, exchanges may wish to submit higher bids than what is needed to win in order to increase their average historical bid and move earlier in the waterfall.
[89] The waterfall can be analyzed in comparison to the other common auction formats. The closest standard auction format to waterfalling is a posted-price mechanism. In a posted-price mechanism, exchanges would be visited one at a time, offered the impression at a personalized reserve, and could either pay the reserve to receive the impression, or pass. In a waterfall, the only difference is that exchanges can choose to pay above the reserve.
[90] In general, under posted-price mechanisms, the first bidder that clears their reserve is awarded the auctioned item.
[91] *See, e.g.*, Hartline and Roughgarden. "Simple versus Optimal Mechanisms." *Proceedings of the 10th ACM Conference on Electronic Commerce*. 2009. pg. 225-234; Alaei, Hartline, Niazadeh, Pountourakis, and Yuan. "Optimal auctions vs. anonymous pricing." *Games and Economic Behavior* vol. 118. 2019. pg. 494-510; Jin, Lu, Qi, Tang, and Xiao. "Tight Approximation Ratio of Anonymous Pricing." *Proceedings of the 51st Annual ACM SIGACT Symposium on the Theory of Computing*. 2019. pg. 674-685.
[92] *See, e.g.*, Sleator and Tarjan. "Amortized efficiency of list update and paging rules." *Communications of the ACM* vol. 28, no. 2. 1985. pg. 202-208; Krengel and Sucheston. "On semiamarts, amarts, and processes with finite value." *Probability on Banach spaces*. 1978. pg. 197-266; Martin L. Weitzman. "Optimal Search for the Best Alternative." *Econometrica* vol. 47, no. 3. 1979. pg. 641–654.

A. Dynamic Allocation

103. The impact of Dynamic Allocation on the waterfall process depends on the types of line items present in the waterfall. More specifically, whether there are only static demand sources or there are both static and live demand sources[109] affects how Dynamic Allocation works, as well as its impact on the auction procedure and outcomes. Hence, I analyze Dynamic Allocation separately for static demand sources and live demand sources.

1) Dynamic Allocation with static demand sources

104. I first present an overview of Dynamic Allocation during the period when it was first introduced.[110] Initially, all line items were static, so Dynamic Allocation addressed a natural shortcoming of the waterfall format. When all line items competing with AdX are static, **Dynamic Allocation with Static Line Items** adjusts the waterfall process in the following manner:

a. First, Google's ad server DFP processes the high priority line items[111] that are not affected by Dynamic Allocation (such as direct deals). If any high priority line item succeeds, the impression is sold, and the waterfall terminates without continuing to subsequent steps.

b. Every low priority line item, including AdX, has both a price floor and a Value CPM.[112] Next, DFP selects the highest Value CPM among all low priority static line

[109] Throughout this report, I use the term "static line item" when referring to line items that do not correspond to outcomes of any auctions. For example, the line items that Google documentation refers to as sponsorship or standard would be static line items. *See* Google. "Line item types and priorities." Accessed on May 31, 2024. https://web.archive.org/web/20240216154938/https:/support.google.com/admanager/answer/177279?hl=en
I use the term "live demand sources" when referring to the ad exchange line items. They are "live" since they hold an auction before submitting a clearing price.

[110] Dynamic Allocation was introduced by DoubleClick, prior to Google's purchase of the company. DoubleClick documentation from that time points to 2007 as the introduction of Dynamic Allocation. *See* Google. "DoubleClick Advertising Exchange." Accessed on May 31, 2024. https://web.archive.org/web/20071001100309/http:/www.doubleclick.com/products/advertisingexchange/index.aspx
Google documentation claims it is 2008, while agreeing that it predates Google's acquisition of DoubleClick. GOOG-AT-MDL-008991406 at -6. ("2008—Dynamic Allocation […] pre doubleclick acquisition.")

[111] Throughout the report, I use "guaranteed" and "high priority" interchangeably when referring to line items that are at priority 1-10 by Google's standards. Similarly, I use "non-guaranteed" and "low priority" interchangeably when referring to line items that are at priority 12-16. *See* Google. "Line item types and priorities." Accessed on May 31, 2024. https://web.archive.org/web/20240216154938/https:/support.google.com/admanager/answer/177279?hl=en
This is in line with the terminology that Google uses. A Google engineer stated that "On a particular request, a standard or sponsorship line item was determined to be guaranteed if the publisher has configured it with a higher priority than AdX, AdSense, and any remnant line item eligible on that request." GOOG-AT-MDL-008842393 at -97. August 4, 2023. "Declaration of Nitish Korula."

[112] Value CPMs are set by the publishers, and they usually correspond to the value of those line items for the publishers. Google provides the following formula to estimate the value CPM: Value CPM = (Total revenue received from ad tags associated with selected line item/Total number of impressions Ad Manager sent to the selected line

108. To illustrate how Dynamic Allocation with Static Line Items works, imagine an impression arrives for a user over the age of 25 who likes running. DFP notices that this satisfies the coarse targeting criteria for a direct deal with Altra, displays Altra's ad, and the waterfall (even with Dynamic Allocation) ends here. This example is illustrated in Figure 18 below.

Figure 18: An impression is allocated to an Altra direct deal in Dynamic Allocation because it fulfills the targeting criteria

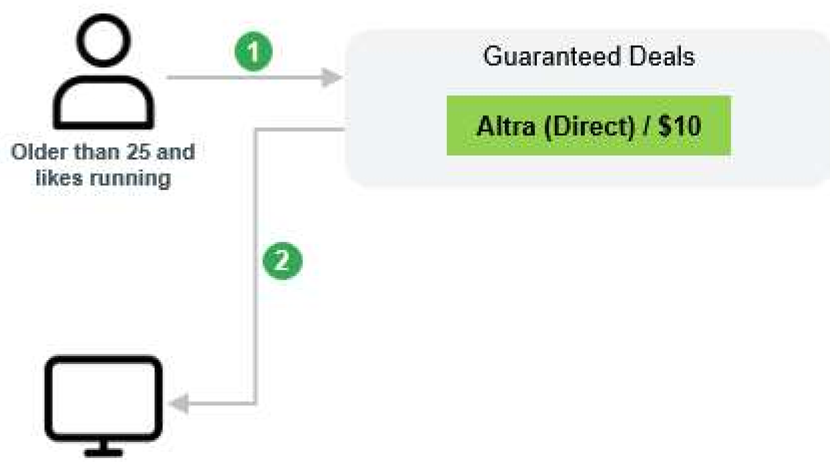

109. Another new impression arrives for a user over the age of 25 who likes expensive handbags. DFP does not have a direct deal that meets this targeting criteria, and so moves on to lower-priority line items and observes that the highest Value CPM option is from a static line item for $4. DFP then calls AdX with a reserve of $4. AdX does not find a buyer above $4, and the waterfall concludes by allocating the impression to the static line item for $4. This example is illustrated in Figure 19 below.

138. Enhanced Dynamic Allocation expands Dynamic Allocation to also include high-priority line items. In particular, the 'Dynamic Allocation Framework' can be interpreted as (a) finding a set of line items, then (b) allowing AdX to be considered before any of those line items,[182] but (c) setting the reserve price of AdX to be equal to the maximum Value/Temporary CPM of those line items.[183]

139. Under Enhanced Dynamic Allocation, impressions that otherwise would have been reserved for high priority line items such as direct deals are instead available for AdX's auction. This follows immediately from the definition of Enhanced Dynamic Allocation since it allows AdX to run an auction for all impressions. The Enhanced Dynamic Allocation-generated reserve price might be high, but AdX will always have the opportunity to run an auction for all impressions. Without Enhanced Dynamic Allocation, any impression with a viable high priority line item would not be available to AdX for auction.

140. Furthermore, AdX is the only exchange that unconditionally has this opportunity. Under Enhanced Dynamic Allocation, another exchange *can* have this opportunity, but only if (a) its Value CPM exceeds the highest temporary CPM among high priority line items, and (b) AdX fails to clear its reserve. In particular, (a) suggests a high barrier to this exchange being considered in front of the high priority line item at all,[184] and (b) notes that AdX still gets a first bite, even if the Value CPM of an exchange is high enough to satisfy (a).

141. As Enhanced Dynamic Allocation runs live auctions for every impression, it will likely create a revenue increase for the publishers in the short run.[185, 186] This conclusion is supported by an internal Google slide deck,[187] an excerpt from which is reproduced in Figure 28. It shows that Enhanced Dynamic Allocation led to an increase in publisher revenue from AdX in the first

---

[182] When one of these line items is a header bidding line item, Dynamic Allocation considers AdX before that line item, although of course that line item itself was generated from a live bid in an auction executed before the ad server.
[183] Under default publisher behavior. As previously noted, a sophisticated publisher could further increase AdX's reserve beyond the maximum Value CPM.
[184] Recall that "On average, direct ads sell for two to four times higher than programmatic ads. Direct ads average $10-20 CPMs (cost per thousand impressions), while programmatic ads average $1-5 CPMs." Google. "Understand Direct and Programmatic Ad Revenue." Accessed on May 31, 2024.
https://web.archive.org/web/20231226200704/https:/newsinitiative.withgoogle.com/resources/trainings/grow-digital-ad-revenue/understand-direct-and-programmatic-ad-revenue/
This suggests that temporary CPMs of high priority line items would likely be higher than static Value CPMs of low priority line items.
[185] In my opinion, the magnitude of these changes would increase due to conduct such as Dynamic Revenue Sharing which causes AdX to clear its publisher-set price floor more often.
[186] Importantly, it is possible that many publishers deprioritize short term revenue, and care more about their revenue in the long run. I elaborate on the long run effects, namely the "cream-skimming effect," below.
[187] GOOG-NE-03872763. "Discussion on improving AdX & AdSense backfill."

quarter of 2014. This also makes sense, as high priority line items are all static. I previously noted in Section IV.A that it would increase revenue to allow a live demand source the option to outbid a static demand source.

Figure 28: An excerpt from an internal Google slide deck plotting the impact of Enhanced Dynamic Allocation on publisher metrics[188]

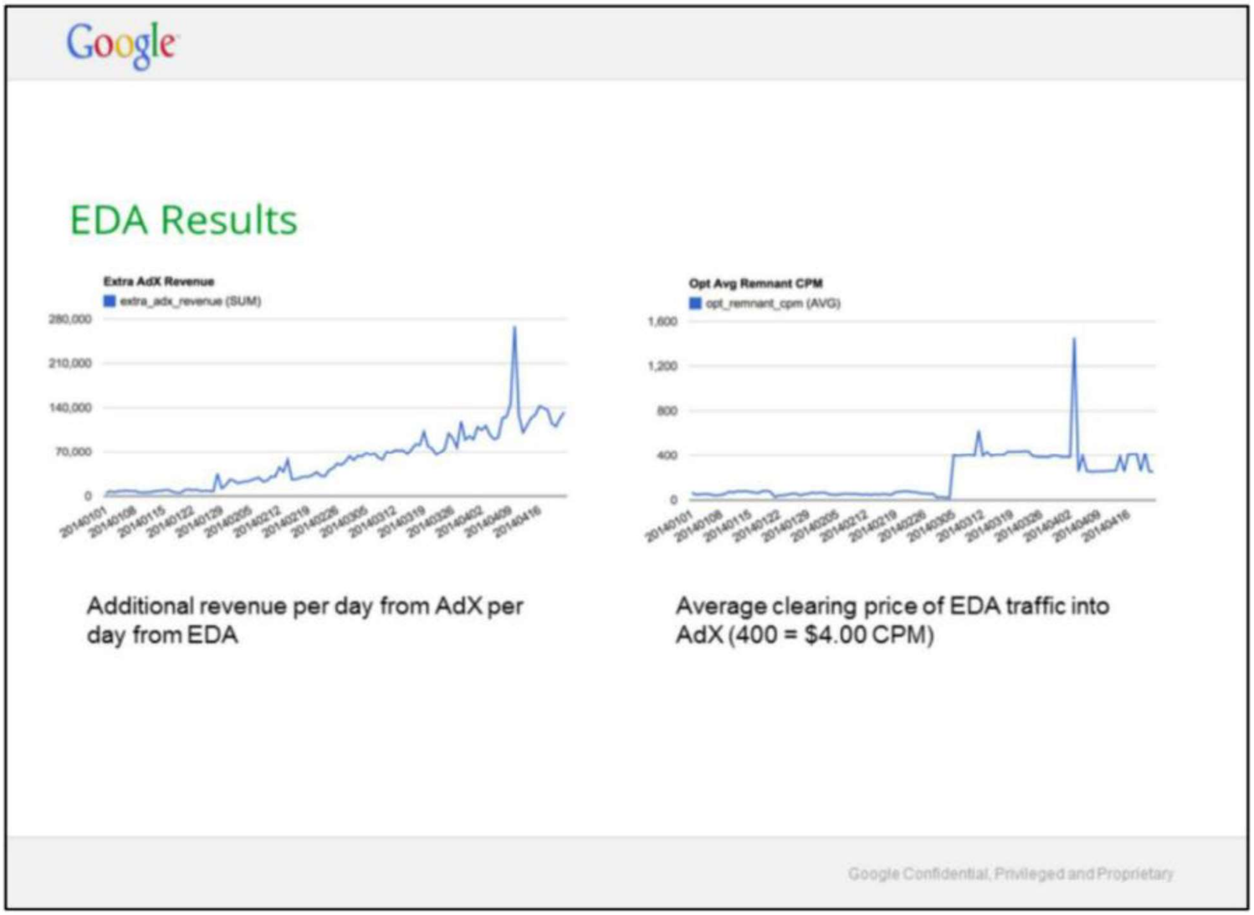

142. Since AdX generates revenue by taking a cut of the clearing prices, the plot above shows that Enhanced Dynamic Allocation increases AdX and Google revenue as well. This can be seen by the upward trends of the plots in the figure above. If impressions that satisfy targeting criteria for direct deals are on average more valuable than impressions that do not,[189] then Enhanced

[188] GOOG-NE-03872763 at -85. "Discussion on improving AdX & AdSense backfill."
[189] Google's online documentation states that "Direct ads average $10-20 CPMs (cost per thousand impressions), while programmatic ads average $1-5 CPMs." *See* Google. "Understand Direct and Programmatic Ad Revenue." Accessed on May 31, 2024. https://web.archive.org/web/20231226200704/https:/newsinitiative.withgoogle.com/resources/trainings/grow-digital-ad-revenue/understand-direct-and-programmatic-ad-revenue/

A. Exchange Bidding

146. I first outline Google's rival product to header bidding. In 2018, Google released **Exchange Bidding**, a technology similar to header bidding, after observing the rise of header bidding technology.[192, 193, 194] Exchange Bidding created a separate auction of auctions into which all exchanges except for AdX submit their bids. AdX would then submit a bid against the prevailing bid from this auction of auctions. Similar to header bidding, Exchange Bidding was a first-price auction.[195] Furthermore, publishers could use both header bidding and Exchange Bidding at the same time.

147. Exchange Bidding augments the waterfall process in the following manner:[196]

a. [redacted]

b. [redacted]

[192] First the tool was called "Exchange Bidding Dynamic Allocation," later renamed to "Exchange Bidding," then to "Open Bidding" after the co-rollouts of Unified Pricing Rules and the AdX auction format change to the first-price. *See* AdExchanger. "Google's Exchange Bidding Is Now 'Open Bidding'; Market Researchers Slip" (August 27, 2019). Accessed on May 31, 2024. https://web.archive.org/web/20220523024855/https:/www.adexchanger.com/ad-exchange-news/tuesday-27082019/
The internal development code name for the product was "Jedi." *See* GOOG-NE-03995243 at -3. July 25, 2018. "PRD: Unified 1P auction and Pricing rules."
[193] To be more precise, the best product to compare to Exchange Bidding is the header bidding variant called "server-side header bidding," since both handle the auction on a server rather than the user's internet browser. *See* Anthony Vargas. "AdExplainer: Client-Side vs. Server-Side Header Bidding: What's The Difference?" (December 1, 2023). Accessed on May 31, 2024. https://web.archive.org/web/20240314163210/https:/www.adexchanger.com/adexplainer/adexplainer-client-side-vs-server-side-header-bidding-whats-the-difference/ ("With client-side header bidding, the bulk of that processing occurs on the user's device in the web browser itself. With server-side header bidding, the processing happens on a remote server.")
[194] Google's First Am. Resps. and Objs. to Plaintiff's Third Set of Interrogs. (May 24, 2024) at 11.
[195] Internal Google documents state that "EB was our first attempt at running a 1P [first-price] auction; Since other exchanges already have experience with submitting 1P bids into HB wrappers, it was the easiest way to build out the product." GOOG-NE-13494966 at -71. May 2019. "Managing Yield."
[196] GOOG-TEX-00000744. April 26, 2017. "Exchange Bidding (Jedi) Open Beta Sates Readiness Review." (internal Google slide deck that discusses how Exchange Bidding worked during its initial implementation, as well as Google's plans on Exchange Bidding rollout.)

advertisers who bid the same whether or not AdX has a Last Look advantage.[224, 225]

160. In addition, when AdX wins these additional impressions, in many cases it pays just a penny more than the winning header bid and so does not increase publisher revenue. Specifically, in cases where AdX's highest value v exceeds the highest header bid h, but h exceeds both AdX's second highest value and its reserve, then AdX wins the impression with Last Look at price h (plus a penny), while without Last Look the header bid of h would have won. When some exchanges participate in header bidding while others participate in exchange bidding,[226] only those who participate in header bidding are vulnerable to AdX's Last Look advantage.[227] Specifically, the highest header bidding bid would become the reserve for AdX, whereas bids through Exchange Bidding are not revealed to AdX before AdX submits its own bid. That is, within AdX and Exchange Bidding exchanges, no one has a Last Look Advantage over the other, because their bids are submitted simultaneously.[228] Exchanges that are integrated into Exchange Bidding see the same DFP reserve as AdX, hence they also have a Last Look advantage over header bidding exchanges.[229]

161. Therefore, one interpretation of Exchange Bidding is that it creates two tiers: Exchanges that participate in header bidding and exchanges that participate in Exchange Bidding together with AdX.[230] Exchanges that participate in header bidding submit bids without seeing others' bids, and therefore have no Last Look advantage over anyone, and are vulnerable to AdX's and Exchange Bidding's Last Look advantage (placing them in the lowest tier).[231] AdX and Exchange Bidding exchanges have a Last Look advantage over exchanges that participate in header bidding, but do not have a Last Look advantage over each other, placing them in the top tier.

[224] Recall again that I have shown a natural example in Appendix D where bidding identically in these situations is optimal for advertisers.
[225] If a sophisticated publisher mildly inflates AdX's reserve specifically because of the Last Look advantage, or if advertisers bid similarly in these two cases, the same conclusions still qualitatively hold. If a sophisticated publisher drastically inflates AdX's reserve specifically because of the Last Look advantage, or advertisers drastically change their bids specifically due to AdX's Last Look advantage, the impact is less clear-cut and would require a complicated analysis weighing the benefits of Last Look versus the impact of an increased reserve and distinct bids.
[226] Importantly, exchanges could participate in both. I am not claiming that this was the case, I am providing explanations based on a hypothetical.
[227] Last Look was phased out in 2019, when AdX transitioned to a first-price auction format. *See* GOOG-TEX-00841386 at -88. "Adx First Price Auction." ("removing last look.")
[228] *See* GOOG-TEX-00000744 at -54. April 26, 2017. "Exchange Bidding (Jed') Open Beta Sates Readiness Review." (diagram shows the '3p floor' entering AdX, but not other exchanges.)
[229] GOOG-DOJ-AT-01815211 at -222. October 2019. "Open Bidding (fka Exchange Bidding)Training."
[230] Again, exchanges can happen to be in multiple of these groups, since they can integrate into both header bidding and Exchange Bidding.
[231] As previously noted, if exchanges that are integrated into header bidding see the same DFP reserve as AdX, header bidders are also vulnerable to a Last Look from these exchanges.

Because a Last Look advantage is significant in first-price auctions, it would be natural for exchanges to want to remove the top tier's Last Look advantage over them and to gain a Last Look advantage over header bidders,[232] even though DFP takes a 5% fee on top of the clearing price when the winner is an Exchange Bidding exchange.[233, 234]

VI. CONDUCT ANALYSIS: UNIFIED PRICING RULES

162. In this section, I provide an analysis of Google's Unified Pricing Rules (UPR), which was instated in 2019 (and in place today[235]) along with Google's ad exchange AdX's transition to the first-price auction format.[236]

163. I demonstrate that UPR leads to lower revenue for the publishers. I also demonstrate that UPR can lead to better win rate and revenue for Google's ad exchange AdX as well as for Google's ad buying tools and lower the win rate and revenue for rival exchanges and ad buying tools.

164. Prior to UPR, publishers could set different reserves that applied to different exchanges or different ad buying tools. Under UPR, publishers can no longer employ these personalized reserves to their full extent,[237] because any reserve price set for non-guaranteed line items applies to all non-guaranteed line items.[238] This reduces publisher choice by preventing them from setting personalized reserve prices. Publishers retain the ability to set personalized reserves on individual advertisers, but not on individual exchanges or ad buying tools.[239] Furthermore, publishers may

---

[232] And after the change referenced in GOOG-DOJ-AT-01809483 went live, it would be further natural for exchanges to want a Last Look over header bidders. GOOG-DOJ-AT-01809483 at -89. March 2017. "Exchange Bidding in Dynamic Allocation (fka Project Jedi)."
[233] Of course, the most natural auction format is to avoid creating tiers and a Last Look advantage at all, and to simply have all exchanges submit bids simultaneously without seeing each other's.
[234] Last Look advantage was removed in 2019 during the implementation of Unified Pricing Rules and AdX's switch to the first-price auction format. *See* GOOG-TEX-00841386 at -89. "Adx First Price Auction."
[235] Google. "Unified pricing rules." Accessed on May 31, 2024. https://web.archive.org/web/20230208153751/https:/support.google.com/admanager/answer/9298008?hl=en (current Google Ad Manager documentation on UPR).
[236] *See generally* Jason Bigler. "An update on first price auctions for Google Ad Manager" (May 10, 2019). Accessed on May 31, 2024. https://web.archive.org/web/20240122142404/https:/blog.google/products/admanager/update-first-price-auctions-google-ad-manager/
[237] GOOG-AT-MDL-000875073 at -83. August 2019. "The Unified First Price Auction."
[238] *See* Google. "Unified pricing rules." Accessed on May 31, 2024. https://web.archive.org/web/20230208153751/https:/support.google.com/admanager/answer/9298008?hl=en (current Google Ad Manager documentation on UPR).
[239] More specifically, publishers are free to set their reserve price at any level they desire, however this reserve price applies to all exchanges and ad buying tools.

they submit a bid, b, between the reserve, *r*, and their value, v, they will win the impression and pay b < v which is strictly better than bidding their true value.

A. Dynamic Revenue Sharing v1

189. DRS was launched in August 2015[278] without announcing it to publishers or advertisers.[279] In its first iteration, DRSv1, AdX dynamically decreased its take rate to be lower than 20% to win impressions that it would not have if the take rate was kept at 20%. AdX decided to either impose a 20% take rate or decrease the take rate depending on a few factors including comparisons between the first and the second highest bid and the publisher reserve, [280] as well as the average take rate among the auctions for that publisher's impressions in that billing period.[281]

190. [redacted]:[282]

a. [redacted]

[redacted] ≥ [redacted]

[redacted] ≥ [redacted]

[redacted] ≥ [redacted]

[278] GOOG-TEX-00777528 at -30. Email thread, "Subject: Re: [Monetization-pm] Re: [drx-pm] LAUNCHED! AdX Dynamic Revenue Share (DRS)." (The email from September 2nd, 2015 states that "Last week we launched Dynamic sell-side Revenue Share (DRS).")
[redacted]
[280] GOOG-NE-06864639 at -43. May 9, 2014. "Dynamic Sell-side Revshare on AdX." (under the subsection "Cases.")
[redacted]
[282] This conduct occurred exclusively during a period where AdX ran a second-price auction.
[283] For the sake of clarity, I abstract away from the ad buying tool fee, since it is immaterial to the analysis conducted here.

i. [redacted]

[redacted]

e. [redacted]

191. DRSv1 impacts the auction only in one case where the highest bid is high enough to clear the reserve price, but not high enough to do so if AdX takes its full fee of 20% of the clearing price. When that happens, AdX dynamically decides to decrease its take rate so that it returns a successful bid to the ad server and wins the impression.

192. To illustrate how DRSv1 works, imagine an impression arrives, and the publisher reports a price floor of \$10.[288] AdX solicits bids and receives top-two bids of \$20 and \$10. In this case, with both a regular second-price auction and DRSv1, the AdX clearing price is \$12.5, because \$20 ≥ \$12.5 > \$10. AdX takes a 20% take rate of \$2.5 and passes on \$10 to the publisher and DRSv1 has no impact. This example is illustrated in Figure 32 below.

[redacted]

[288] For ease of exposition, I use 'price floor' for when the publisher sets a reserve on an exchange, and 'reserve' for when an exchange sets a reserve on advertisers.

between the first and the second highest bid and the publisher reserve, as well as debt balances for the publishers and advertisers.[294]

197. DRSv2 was launched in the second half of 2016.[295] Google announced DRSv2 when it was launched.[296] The publishers were allowed to opt out of DRSv2, however, if they did, Google turned off DRSv1 for these publishers as well.[297] Advertisers and ad buying tools could not opt out of DRSv2.[298]

198. [redacted]

a. [redacted][299]

b. [redacted] ≥ [redacted]

c. [redacted] ≥ [redacted]

d. [redacted] ≥ [redacted]

---

[294] [redacted] 2. [redacted]

[295] GOOG-NE-04934281 at -81. July 30, 2018. "Dynamic Revenue Share." (describing release dates of "feature flags" (functionality), including that it launched into AdX UI in June 2016 and came into effect in August 2016.)

[296] GOOG-NE-06842715 at -20. May 10, 2016. "AdX Auction Optimizations." (describing that DRS would be announced in June 2016.)

[297] GOOG-NE-04934281 at -86. July 30, 2018. "Dynamic Revenue Share." ("You may choose to opt-out of revenue share based optimizations in the AdX UI. If you opt-out we will apply your contracted revenue share to every Open Auction query and you will not benefit from the increased revenue from this optimization.")

[298] GOOG-NE-04934281 at -85. July 30, 2018. "Dynamic Revenue Share." ("Q: Can buyers opt out? A: Revenue share based optimizations are controlled by sellers only.")

[299] For the sake of clarity, I abstract away from the ad buying tool fee, since it is immaterial to the analysis conducted here.

[300] [redacted]

207. I discuss the implications of participating in a debt-aware second-price auction later in this section, and observe now just that step d.ii. above indeed charges the winning bidder a price that exceeds their bid.[311]

C. Truthful Dynamic Revenue Sharing

208. With the third and last iteration tDRS, AdX dynamically adjusted its take rate to sometimes be higher or lower than 20% to win impressions that it would not have if the take rate was kept at 20% on a per-query basis. Under tDRS, Google determined the dynamic take rate it is going to charge before 'peeking' at the bids.[312] In contrast, both DRSv1 and DRSv2 adjusted the take rate after AdX observed the submitted bids.[313] Under tDRS, the take rate calculation is done based on the past AdX data. Internal Google documents states that the prediction model that determines the take rate "predicts for a given query whether a specific buyer would bid above the pre-revshare reserve price."[314]

209. tDRS was fully launched in the second half of 2018.[315] When AdX migrated to a first-price auction format in 2019, the DRS program was shut off.[316]

210. [redacted]

a. [redacted][317]

b. [redacted]

[311] And recall again that this follows because we are assuming all debt clears, and so therefore can count payments made to clear debt the moment that debt is accumulated.
[312] GOOG-AT-MDL-019244499. "Truthful DRS Auction Walkthrough." ("For each buyer, its reserve price revshare factor will be determined based prediction result before the request is being passed down to RTBs or CAT2 mixer (for Adwords and DBM).")
[313] GOOG-NE-13226622 at -2. "Truthful DRS Design Doc." ("One known issue with the current DRS is that it makes the auction untruthful as we determine the AdX revshare after seeing buyers' bids and use winner's bid to price itself (first-pricing) when the bid is within the dynamic region.")
[314] GOOG-NE-13214748 at -8. "Modeling Design Doc for Truthful DRS."
[315] GOOG-TEX-00858434. January 29, 2020. "Dynamic Revenue Share." ("Update (July 30, 2018): We launched a new DRS model (tDRS).")
[redacted]

2) Advertisers would have submitted different bids to maximize their payoffs had Google revealed DRSv1

227. Google misled advertisers by not revealing DRSv1, and hence led them to believe the AdX auction was a regular second-price auction, which would cause them to engage in suboptimal behavior. When advertisers believe they are participating in a regular second-price auction, they would bid their true value for the impression, because it is a truthful auction. However, DRSv1 is not truthful, as established before. Therefore, concealing DRSv1 caused advertisers to bid their true value in a non-truthful auction, whereas advertisers would get higher a higher gain by shading their bids.

228. By not revealing DRSv1 to the advertisers, Google made material gains. This is because if advertisers were to shade their bids, which is the natural bidding behavior in a non-truthful auction like DRSv1, this would lead to less revenue for both AdX and publishers.[347] However, advertisers likely did not shade their bids, since Google never publicly revealed DRSv1.

G. Some aspects of DRS are exceptionally misleading

229. To conclude the section, I want to briefly note a few aspects of DRSv2 that I find exceptionally misleading to advertisers, and an aspect of DRSv2 and tDRS that I find misleading to publishers. Much of my analysis below concerns the concept of 'debt' to mislead both advertisers and publishers regarding how much they are paying or paid out.

230. First, I want to repeat that my previous analysis establishes that when advertisers behave optimally in DRSv2, no transactions should ever occur in the dynamic region. Instead, Google claims that enough transactions occurred in the dynamic region to account for a [redacted][348] This increase necessarily comes at the expense of advertisers ultimately *paying more than their value for an impression*.

231. Next, I want to highlight aspects of Google's description[349] regarding DRS that I find misleading, due to the concept of debt.

---

[347] But this does not mean DRSv1 as a whole, after accounting for bid shading, would necessarily lose revenue. It merely means that when comparing "DRSv1 when advertisers shade their bids" yields higher payoff for advertisers and less revenue for AdX and the publishers than "DRSv1 when advertisers bid their true values," which motivates concealing DRSv1.
[348] GOOG-NE-13234466 at -67. "Overall Pub Yield With DRS(v2)."
[349] GOOG-NE-04934281 at -84. July 30, 2018. "Dynamic Revenue Share."

a. For DRSv2, Google states: "Buyers are never charged more than their bid," I find this claim exceptionally misleading to advertisers. If by 'charged', one means 'charged on this impression as immediate payment, ignoring any debt that will be paid later', then the sentence is technically true. But if by 'charged', one means 'charged on this impression either as immediate payment or as debt to be collected later', then the sentence is false. Indeed, any winner in the dynamic region is ultimately charged more than their bid after accounting for both immediate payment and debt to be collected later.

b. For DRSv2, Google states: "sellers are always paid at least their reserve." I find this claim misleading to publishers. If by 'paid', one means 'paid on this impression as immediate payment, ignoring any debt generated that must be paid back later', then the sentence is technically true. But if by 'paid', one means 'paid on this impression as immediate payment, after subtracting any debt assigned that will be collected later', then the sentence is sometimes false. On impressions transacted in the dynamic region, publishers are indeed paid their reserve as immediate payment, but are also assigned non-negative debt. This debt may wind up being cancelled (if the winning advertiser clears their own debt while later transacting with this publisher), owed (if the winning advertiser clears their own debt elsewhere), or yielding extra payout (if some other advertiser later clears debt incurred elsewhere with this publisher). Even after accounting for debt, it is plausible to say that "sellers as a whole are paid at least the reserve set on any cleared impression." But it is inaccurate to say "sellers are *always* paid at least their reserve," because some sellers are not ultimately paid their reserve.[350]

c. If the concept of debt was not clearly disclosed, the general description of DRS as per-query revenue share optimization is insufficient for advertisers to draw conclusions at the level I have drawn in my report. Moreover, even for advertisers who are already optimizing bids at a per-impression level, the concept of debt significantly obscures feedback. Indeed, a typical optimizer might ask questions of the form "if I change my bid on this impression, what change does that cause in my payoff from this impression?" The concept of debt now means that changing a

[350] Sellers would reasonably care about whether they are paid their reserve or not, as this reserve constitutes the minimum amount they've decided to accept in order to forego the opportunity cost of selling the impression elsewhere.

Global Bernanke, without improving GDN buyers' utilities), which leads to an increased win rate and revenue for GDN (again in the case of Projects Bernanke and Global Bernanke, without assisting GDN advertisers at all).

233. Project Bernanke and all its variants can be understood as simultaneously facilitating the effects of collusion among GDN advertisers, without their knowledge, and overbidding in auctions. I explain this view in detail in this section. The starting point for Project Bernanke is Google's observation that [redacted][354] Because GDN is 'second-pricing itself', GDN would benefit by lowering the second-highest bid it sends in order to lower the payment GDN must make to win the impression.[355] In isolation, this would be a pure transfer of funds from exchange/publisher to GDN, but would technically result in a high take rate by GDN towards its advertisers, compared to what is contracted. This aspect of Project Bernanke is akin to facilitating collusion among GDN advertisers (and in this case, without their knowledge).[356] The second half of Project Bernanke uses the savings from the first half and spends it to subsidize overbidding.[357] That is, the second half of Project Bernanke boosts the bids of its advertisers before sending them to AdX, but uses the funds from the first half to cover any payment made above the advertiser's true bid. In all Project Bernanke variants, the two halves balance out to (a) generate increased revenue and increased win rate for GDN, (b) balance GDN's take rate at the intended 14%,[358] (c) have an indeterminate effect on publisher revenue (the first half decreases publisher revenue while the second half increases it), (d) in the case of Project Bernanke and Project Global Bernanke, not improve GDN advertisers' payoffs at all. Below I describe Project Bernanke in greater detail, and Appendix H contains additional discussion on overbidding and collusion in first- and second-price auctions.

[redacted]

356 By using the word 'collusion', I do not mean to imply that it is 'wrong' from a pure auction theory perspective for a group of bidders to get together and jointly strategize on how to collectively bid, nor for an ad buying tool to facilitate this. My understanding is that other ad buying tools may have dropped their second-highest bid entirely. GOOG-NE-13200831 at -1. "The case for encouraging buyers to declare two bids." ("Currently, the only [AdX] buyer who is employing this strategy [sending two bids] is GDN.")
Still, GDN is indeed facilitating collusion by implementing a joint strategy for its advertiser pool together, rather than processing each advertiser's bid in isolation.

357 GOOG-AT-MDL-001412616 at -20. "Project Bernanke and margins story." ("What if overbid? We could bid too much. But we have to subsidize it. One is good for us [GDN] and bad for publishers. Other is bad for us [GDN] and good for publishers.")

358 Some documents I have reviewed states that the GDN take rate is 14% and others state that it was 15%. I use 14% throughout the text, except in the cases where I cite a specific document that states 15%. This difference in the take rate does not have any effect on my conclusions throughout this section.

A. Project Bernanke

234. Prior to any modifications, GDN ran an internal auction (called the "CAT2" auction)[359] with only GDN advertisers and submitted the top two bids from that auction to AdX.[360, 361] Under **Project Bernanke**, between 2013 and 2015,[362] GDN manipulated advertisers' bids before sending them to AdX in the following manner:[363]

a. [redacted]

b. [redacted]

c. On a per-auction basis, GDN under Project Bernanke always paid AdX the AdX clearing price (when a GDN bidder wins). GDN charged its highest bidder in the following manner when GDN wins the AdX auction:

i. Let b_1 denote the highest GDN bid, and b_2 denote the second highest GDN bid. Let also c denote GDN's minimum bid to win on AdX (that is, the maximum of AdX's reserve and all non-GDN bids received by AdX). Then the winning bidder is charged [redacted]

---

[359] Also called the "CAT2 auction." GOOG-NE-11753797 at -37. February 11, 2019. "DVAA Quality, Formats, O&O - Q1 2019 All Hands."
[360] GOOG-NE-06839089 at -94. "Project Bernanke." ("GDN submits two bids into AdX auction…")
[361] During the entire lifetime of Project Bernanke, AdX conducted second-price auctions. This subsection covers the original Project Bernanke, and later subsections cover its variants.
[362] Project Bernanke was launched in 2013. It was in place until Project Global Bernanke was launched in 2015. *See* GOOG-DOJ-28385887 at -93, 94. August 17, 2015. "Beyond Bernanke." ("Bernanke (late 2013)… Global Bernanke (mid-2015)...")
[363] GOOG-AT-MDL-008881638 at -8. October 30, 2014. "Rethinking Bernanke: Grid search to line search." (describing the mathematical formulation of Bernanke.)
[364] The 'clean' approach would be multiplying the bid by 0.86, which corresponds to just the GDN take rate.
[365] To be clear, I am opining that this aspect of Project Bernanke, and collusion among GDN bidders, helps GDN at the expense of publishers. This is not an opinion on Project Bernanke as a whole.
[366] For comparison, the 'clean' approach would be to forward 0.86 times the highest bid. Going from 0.86 to 1 can be interpreted similarly to DRS – it guarantees that a GDN Advertiser will win whenever its value exceeds AdX's clearing price, and GDN will lower its take rate. Going from 1 to 4 can be interpreted as overbidding, because if it causes a GDN Advertiser to win when they otherwise wouldn't have, it will necessarily be at a price exceeding their value.
[367] To be clear, I am opining that this aspect of Project Bernanke, and overbidding in a second-price auction, helps publishers at the expense of GDN advertisers – this is not an opinion on Project Bernanke as a whole.

256. Under Projects Bernanke and Global Bernanke, the win rate of non-GDN advertisers on AdX would decrease. This is because the GDN advertisers are still winning every impression that they would have won without Projects Bernanke and Global Bernanke, but they are also winning additional impressions. Some of these impressions previously would have been won by non-GDN advertisers, so these advertisers face a lower win rate.[393]

257. Google internal documentation shows that non-GDN advertisers saw a decline in their win rates. [redacted]

2) Advertisers would have shaded their bids to maximize their payoff had they known about Projects Bernanke and Global Bernanke

258. Google concealed vital information from advertisers by concealing Projects Bernanke and Global Bernanke. Provided that neither Project Bernanke nor Project Global Bernanke were disclosed to advertisers, they would naturally believe they were still participating in a truthful second-price auction and bid their true value as a result. If advertisers knew they were participating in a non-truthful auction, they would have instead considered shading their bids. Knowing the auction format is vital information to advertisers aiming to optimize their payoff. In particular, Projects Bernanke and Global Bernanke are dirty second-price auctions.[397] Specifically, if c denotes the minimum bid to win for GDN on AdX, then from the perspective of a GDN advertiser, Projects Bernanke and Global Bernanke are both dirty second-price auctions with soft floor c and hard floor c/α. That is, as long as the highest GDN bidder exceeds c/α, they will win, because their bid will be increased by α to exceed c. If their bid further exceeds c, the auction turns into a regular second-price auction. If their bid falls between c/α and c, they will pay their bid and be subsidized by Bernanke pool for the remaining amount. Therefore, this is a dirty second-price auction.

[393] Some of these impressions could have been previously unsold.
[394] GOOG-DOJ-28385887 at -95. August 17, 2015. "Beyond Bernanke."
[395] GOOG-DOJ-28386151 at -67. December 10, 2013. "Project Bernanke - Quantitative Easing on the AdExchange."
[396] This is because, as noted previously, the price a non-GDN advertiser pays if they still win under Projects Bernanke and Global Bernanke can only increase from what they would pay with no Bernanke.
[397] See Section VII for more information on dirty second-price auctions.

G. Project First-Price Bernanke

264. AdX eventually switched to a first-price auction,[402] which renders the particular collusion mechanics of old Project Bernanke obsolete. This is because the first-price auction is pay-your-bid, hence dropping GDN's second highest bid does not impact the auction at all. The general framework of colluding and overbidding still apply, but the precise mechanics differ. I provide more details on this in Appendix H.

265. As I previously discussed, first-price auctions are not truthful. In fact, for any bid b less than the bidder's value *v*, it is always better to bid b instead of v.[403] However, an economic approach called the "Revelation Principle"[404] allows an intermediary to make a first-price auction truthful for participants. Intuitively, it works in the following way: The intermediary first comes up with a device that takes in a bidder's value as an input and calculates the optimal bid. The intermediary then tells the bidders to report their true values and assures them that if they win, they will be charged their minimum bid to win. But the intermediary does not submit the true values of the bidders to the auction on their behalf, and instead submits bids calculated by the optimization device. The auctioneer then executes a first-price auction with those bids. From the perspective of the advertisers, this is a truthful auction, because they always pay their minimum bid to win. But potentially there is a mismatch in payments since the minimum bid to win (what is charged to the winning bidder) might differ from what is calculated as the optimal bid from the highest value submitted by the bidders (what is paid to the auctioneer). If the device is excellent at bid optimization, these will perfectly balance out on average. If not, there can be a benefit or loss to either party. For the rest of this analysis, I assume that the bid optimizer is excellent.[405]

266. First-Price Project Bernanke has three components: (a) a bid optimizer for GDN users that makes their participation in AdX's first-price auction truthful, (b) collusion among GDN bidders, which increases GDN's payoff at the expense of publishers' revenue, (c) overbidding, which lowers GDN's and increases publishers' revenue. In comparison to Projects Bernanke and Global

---

[402] *See* Jason Bigler. "An update on first price auctions for Google Ad Manager" (May 10, 2019). Accessed on May 31, 2024. https://web.archive.org/web/20240122142404/https:/blog.google/products/admanager/update-first-price-auctions-google-ad-manager/

[403] To see this, observe that bidding v guarantees a payoff of 0, no matter what since it either leads to losing the item, or winning the item and paying the bid. Bidding b that is less than v instead guarantees a payoff no worse than 0 since it either leads to losing the item or winning the item and paying less than the value.

[404] *See* Roger B. Myerson. "Incentive Compatibility and the Bargaining Problem." *Econometrica* vol. 47, no. 1. 1979. pg. 61–73.

[405] Shortcomings of the bid-optimizer are certainly relevant for thinking through the impacts of First-Price Project Bernanke, but it is not relevant to the conclusions I draw based on collusion and overbidding alone.

Bernanke, (b) and (c) are conceptually similar but implemented via different mechanics (due to the different mechanics between first- and second-price auctions).

267. First-Price Project Bernanke carries the same motivation as Project Bernanke. GDN bidders could collude in AdX's auction to increase GDN's revenue at the expense of the publisher, but this hurts publishers' revenues. First-Price Project Bernanke again observes that overbidding has the opposite effect of lowering GDN's payoff but helping publishers' revenues, although it causes collateral damage to non-GDN advertisers. Additionally, there is also an added complication due to intermediating AdX's first-price auction to make it truthful.

268. Under First-Price Project Bernanke, GDN manipulated advertisers' bids before sending them to AdX in the following manner:[406, 407, 408]

a. [redacted][409]

b. [redacted]

c. [redacted]

i. [redacted]

[406] First Price Project Bernanke was launched in 2019. *See* GOOG-AT-MDL-008842383 at -88. August 5, 2023. "Declaration of Nirmal Jayaram." ("Google updated the Bernanke algorithms in 2019 to be compatible with the Unified First Price Auction. The updated version of Bernanke was sometimes referred to within Google as 'Alchemist.'")
[407] GOOG-DOJ-AT-02224828. March 2019. "The Alchemist." [redacted]
[408] This was also called "The Alchemist." Nomenclature was further subdivided into the "[redacted]" which [redacted]
[409] [redacted]

271. When AdX participates in a simultaneous auction with other exchanges, First-Price Project Bernanke could cause AdX to win more or fewer impressions. When AdX participates in an auction with other exchanges, such as in Exchange Bidding, it matters not only whether AdX clears its reserve, but also at what price it clears. A higher clearing price would cause AdX to win more often, and a lower clearing price would cause AdX to win less often. In a first-price auction, the clearing price is the highest bid, and therefore AdX's clearing price would increase or decrease based on whether Project First-Price Bernanke causes GDN's highest submitted bid to increase or decrease. The overbidding aspect causes GDN to submit a higher bid. On the other hand, the collusion aspect causes GDN to submit a lower bid.[414] Because impacts are possible in both directions, AdX would sometimes have a higher clearing price and sometimes have a lower clearing price.

IX. CONDUCT ANALYSIS: RESERVE PRICE OPTIMIZATION

272. In this section, I provide an analysis of Google's Reserve Price Optimization (RPO)[415] conduct. I demonstrate that RPO leads to higher revenue for Google's ad exchange AdX and explain the mechanisms through which it leads to lower payoff to advertisers and could lead to lower revenue for some publishers. Furthermore, I outline how it impacts publisher and advertiser behavior. The negative effects of RPO to advertiser payoff, and possibly some publishers' revenues, at least partially stem from Google's concealment of the conduct.

273. In particular, it is my opinion that Google concealed information that is material to both publishers and advertisers during the period RPO was concealed. It is also my opinion that even after RPO was revealed, publishers might set suboptimal reserves on any impression for which RPO is a possibility.[416]

274. Under RPO, AdX used data available to them (prior to seeing live bids) to calculate per-buyer reserve prices[417] that it believed would optimize AdX's revenue.[418] AdX then used these

---

[414] Note that the collusion aspect cannot cause GDN to fail to submit a bid above AdX's reserve. But conditioned on meeting AdX's reserve, GDN could submit either a higher or lower bid.
[415] There seems to be other programs that are called RPO previously. However, they substantially differ from the conduct I am discussing here. The main difference between those programs and this conduct is that this conduct sets per-buyer reserve prices.
[416] That is, unless a publisher knows whether RPO activates on a particular impression, and if so what reserve RPO would set, I would expect publishers to lack sufficient information to set a profit-maximizing reserve.
[417] GOOG-NE-13204729 at -36. August 17, 2015. "AdX Dynamic Price." ("Buyer Specific Reserve Prices. Different buyers may get different reserve prices.")
[418] GOOG-NE-06151351 at -52. November 12, 2015. Email thread, "Subject: [Monetization-pm] Re: [drx-pm] LAUNCHED! Dynamic Pricing (RPO) for AdX sellers." ("[RPO] generate[s] a histogram of historical bids and transaction prices […] pick[s] a reserve price that maximizes predicted revenue.")

reserves in its own auction instead of the reserves set by the publisher, although this reserve was always at least as large as the reserve set by the publisher.[419] The program was launched in phases between April and October 2015.[420] Initially, Google did not announce this program to its customers.[421] Later, Google announced the program to its customers under the name "optimized pricing" on May 12th, 2016, more a year after its initial rollout.[422] Publishers were not allowed to opt out of the program.[423] The program was deprecated in 2019 with the switch of AdX to the first-price auction format.[424]

275. Internal Google documents suggests that RPO relies on an algorithmic optimization that "set[s] optimized reserve prices in AdX auction[s]" to "increase the revenue for publishers" via "model[ing] effect of various reserve prices" and then "pick[ing] the best one."[425] Importantly, this tool aims to set the reserve price just below what the highest bidder is willing to pay,[426] by coming up with an empirical estimate of this willingness to pay, which was assumed to be equal their bid due to the truthfulness of the second-price auction.[427] An internal Google document states that the goal of RPO was to "select a reserve price as close to the anticipated first price as possible in order to trade buyer for seller surplus."[428] If AdX has sufficient data to form an accurate prediction of the maximum advertiser value v, the optimal reserve price to set is exactly v.

Different Google internal documents outline different strategies AdX used to employ the data they have to best estimate the RPO reserve prices. Which data was used and how data was processed are not relevant to the conclusions I provide below.

[419] GOOG-NE-03640022 at -2. "AdX Managed Reserves." ("Currently RPO can only raise reserve prices".) Notice that optimizing publisher revenue and AdX revenue are equivalent since AdX revenue corresponds to 20% of the publisher revenue.

[420] *See, e.g.,* GOOG-NE-06151351 at -52. November 12, 2015. Email thread, "Subject: [Monetization-pm] Re: [drx-pm] LAUNCHED! Dynamic Pricing (RPO) for AdX sellers." ("Between April and October we launched and improved new systems to dynamically set auction reserve prices for AdX sellers.")

[421] GOOG-NE-09485306 at -432. December 18, 2017. "OLD – New Ad Manager Indirect Notes." ("We are not commercializing this externally for now.")

[422] *See* Jonathan Bellack. "Smarter optimizations to support a healthier programmatic market" (May 12, 2016). Accessed on May 31, 2024. https://web.archive.org/web/20200929015943/https:/blog.google/products/admanager/smarter-optimizations-to-suppor/

[423] GOOG-NE-06842715 at -18. May 10, 2016. "AdX Auction Optimizations." ("No opt-out possible.")

[424] GOOG-AT-MDL-000987708 at -8. April 9, 2021. "PM Perspective on 1P RPO." ("When we transitioned to a 1st price auction and launched unified pricing rules in September 2019, we had to turn off 2P RPO since it was designed to work in a 2nd price auction (duh).")

[425] GOOG-NE-13204729 at -30. August 17, 2015. "AdX Dynamic Price."

[426] When there is sufficient data to predict the highest bidder's willingness to pay exactly, the goal is indeed to set the reserve price just below this. Often there is insufficient data to predict the highest bidder's willingness to pay exactly, and in these cases Reserve Price Optimization instead aims to set the optimal reserve given the information it has.

[427] GOOG-NE-13204729 at -34. August 17, 2015. "AdX Dynamic Price." (the slide titled "How to Guess the Top Bid" explains how Google thought about estimating the highest bid.)

[428] GOOG-NE-03640022 at -2. "AdX Managed Reserves."